

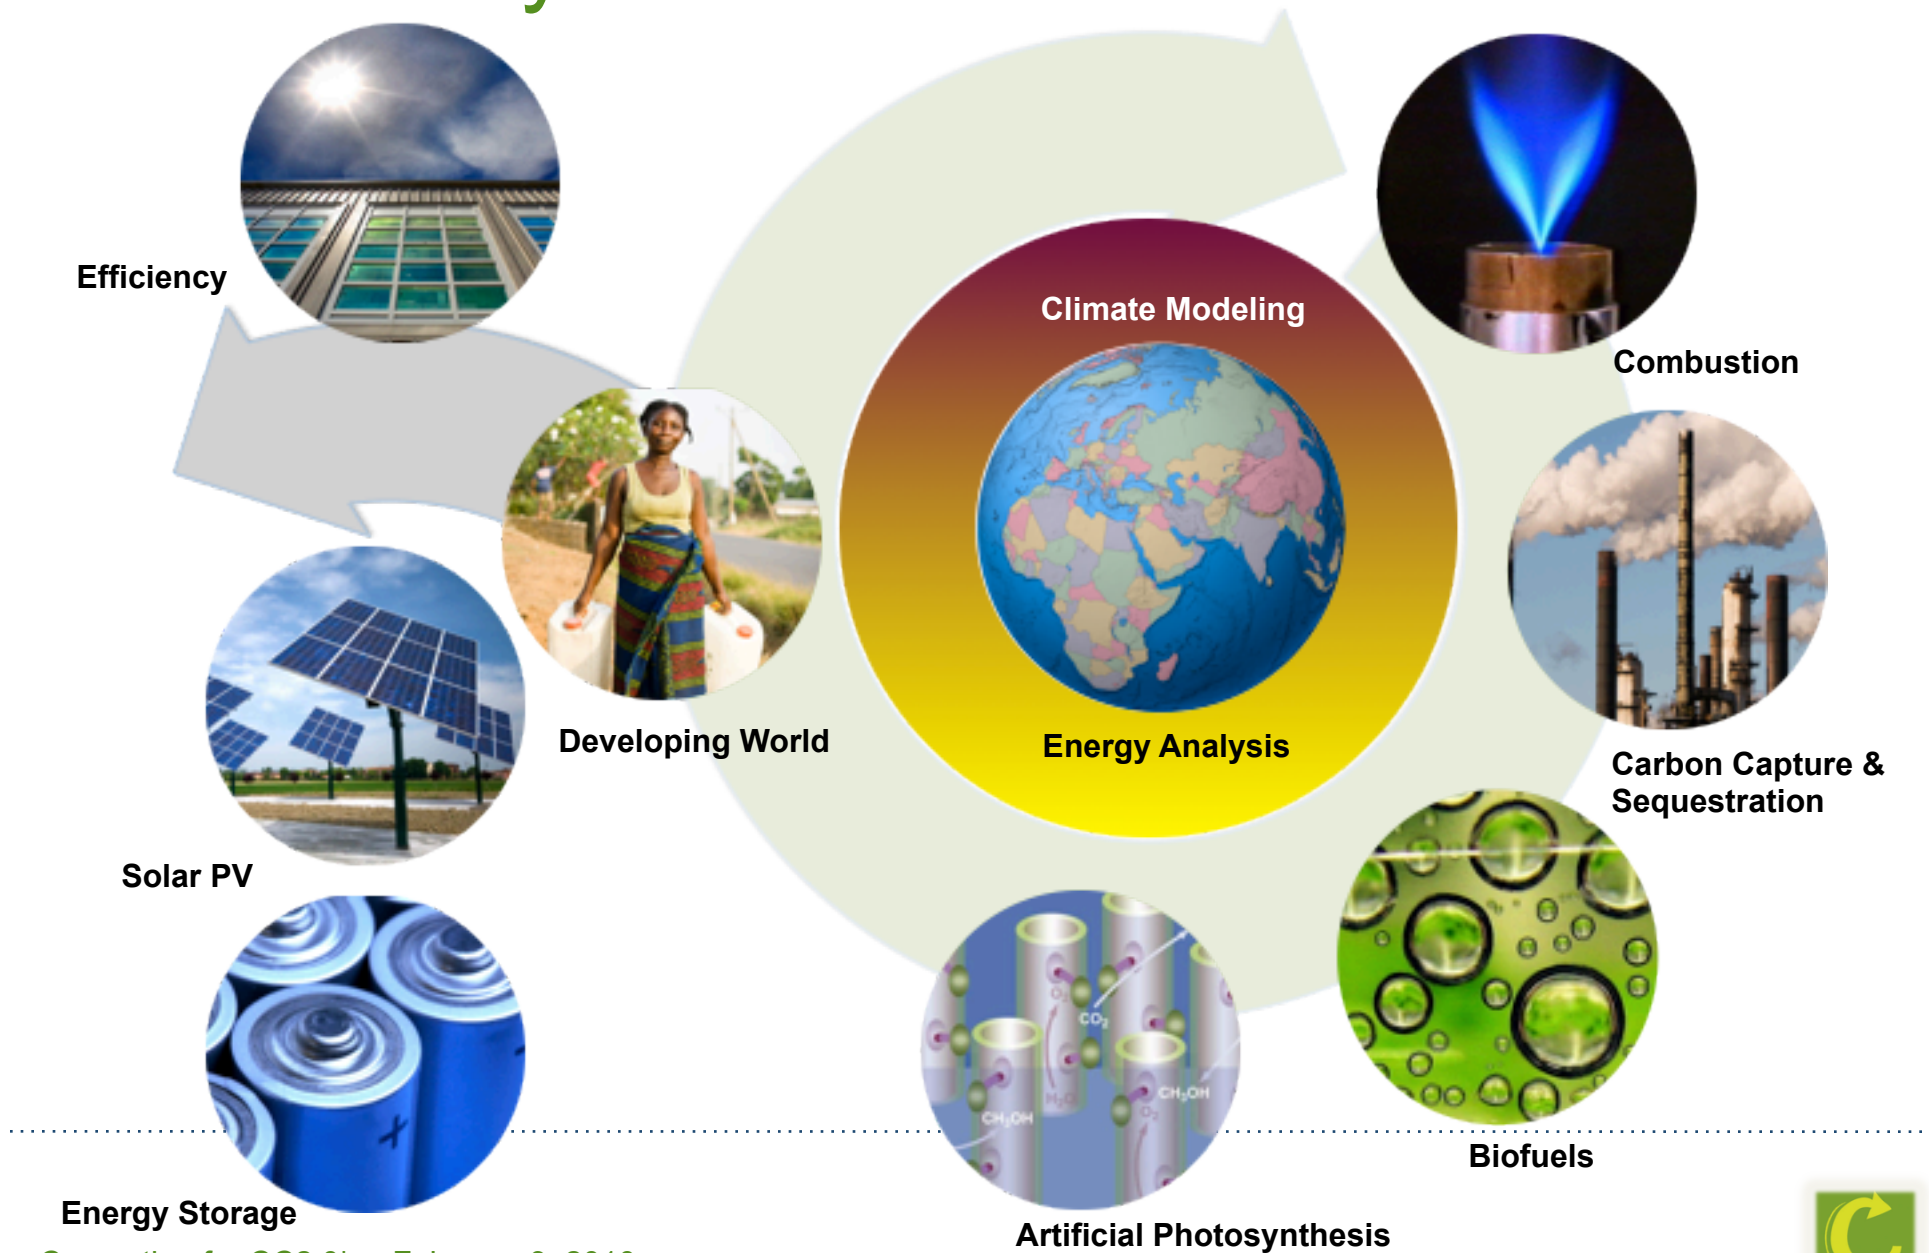
Computation in CC 2.0

Juan Meza

Senior Scientist and Department Head
High Performance Computing Research

February 3, 2010

Carbon Cycle 2.0 Initiative



NERSC 2009 Configuration

Large-Scale Computing System

Franklin (NERSC-5): Cray XT4

- 9,532 compute nodes; 38,128 cores
- ~25 Tflop/s on applications; 356 Tflop/s peak

Hopper (NERSC-6): Cray XT


- Phase 1: Cray XT5, 668 nodes, 5344 cores
- Phase 2: > 1 Pflop/s peak




- **NERSC systems capability in CPU hours**
 - **280M hours in 2010**
 - 28M of that in NERSC Director reserve
 - 28M in ASCR Leadership Computational Challenge (ALCC)
 - **Estimated 1B hours in 2011 (100M in each reserve)**
- **Personnel:**
 - **Computational Science and Engineering Petascale Initiative Program (emphasis on EFRCs)**

The Computational Research Division

Energy technology




Nano systems




Renewable energy


engages in computational science partnerships, developing algorithms, tools and techniques that enable advanced computational modeling and simulation, and lead to new understanding in areas such as



global climate



Combustion processes




Biological systems



Combustion

Simulation of lean premixed hydrogen flames stabilized on a low-swirl burner

- Low Mach number formulation exploits mathematical structure of the problem
 - Advanced numerical methodology, including projection methodology, adaptive mesh refinement, and parallel implementation using BoxLib
 - Detailed chemistry and transport
 - No explicit models for turbulence or turbulence / chemistry interaction
 - 25 cm x 25 cm x 25 cm
- Combined methodology enables simulation at effective resolution of 8B cells (2048^3)
- Simulation captures cellular structure of thermodynamically unstable lean hydrogen flames and provides insight into experimental diagnostics




PI: J. Bell, LBNL

Simulations performed at NERSC under an INCITE grant



Experimental comparison of OH Radical

Simultaneous rendering of OH and Vorticity






Carbon Capture & Sequestration



Knowledge-guided screening tools for identification of new materials for CO₂ separation

Goal: Screening of large databases of porous materials (e.g. zeolites) to identify structures with optimal performance for CO₂ separation.

Computational Challenges: Develop mathematical techniques to characterize and represent the geometrical and physical features of porous structures and to use statistical and data mining methods to identify similarity between structures



These graphs represent the topology of four zeolite structures obtained with Mapper method. Nodes of each graph are obtained by clustering of atoms of similar density and are colored according to values of density.

Similarity is based on Hausdorff-Gromov distance

M. Haranczyk, J.A. Sethian, J.C. Meza



Collaborators: B. Liu, K. Theisen, B. Smit (UCB), J. Kloke, G. Carlsson (Stanford)

Computing for CC2.0| February 3, 2010



Structure Descriptors, Representations and Similarity Measures

Calculation of accessible volume and transport pathways based on analysis of configuration space of a molecule inside zeolite structure using Fast Marching Method – Robotic Navigation (Sethian)




POC: M. Haranczyk (LBNL)

Front propagation to obtain accessible volume


AMR provides accurate onset times

- Accurate and rigorous determinations of the onset time of convection, and the long-term stabilized CO₂ mass flux.
- Showed that the density-driven convective process significantly increases the rate at which CO₂ is transformed into a negatively buoyant state.
- Determined dependence of the mass flux on the formation properties.
- Proposed a simplified model for integration into a full carbon sequestration simulation.
- 3D simulations use 2048 CPUs for more than 48 hours on Franklin at NERSC.




Evolution of convective fingers with time. The arrows show how the velocity field drives the formation and evolution of the fingers. Joint work with Karsten Pruess at ESD.





Extension to realistic geometries will require high resolution images from ALS



Experimental image of packed sintered glass beads
Courtesy of R. Detwiler, J. Nelson (LLNL)



Embedded boundaries obtained from level set and implicit function representation of low resolution image data on grid; with T. Ligocki





Raw porous media image data with various levels of smoothing; Courtesy of PNNL, A. Tartakovsky; with T. Ligocki




Solar PV

Can one use an intermediate state to improve solar cell efficiency?



Highest O induced state




ZnTe bottom of cond. band state


- Single band material theoretical PV efficiency is 30%
- With an intermediate state, the PV efficiency could be 60%
- One proposed material ZnTe:O
 - Is there really a gap?
 - Is it optically forbidden?
- LS3DF calculation for 3500 atom 3% O alloy [one hour on 17,000 cores of Franklin]
- Yes, there is a gap, and O induced states are very localized.
- Winner of ACM Gordon Bell Award for Algorithm Innovation in 2008

L-W. Wang, B. Lee, Z. Zhao, H. Shan, J. Meza, D. Bailey, E. Strohmaier.
INCITE project, NERSC, NCCS.



LS3DF computations yield dipole moments of nanorods and the effects on electrons



$P = 30.3$ Debye



$P=73.3$ Debye



$\text{Cd}_{714}\text{Se}_{724}$

WZ

- ❖ Equal volume nanorods can have different dipole moments
- ❖ The inequality comes from shape dependent self-screening
- ❖ Dipole moments depend on bulk and surface contributions
- ❖ Dipole moments can significantly change the electron and hole wave functions

System Performance Summary



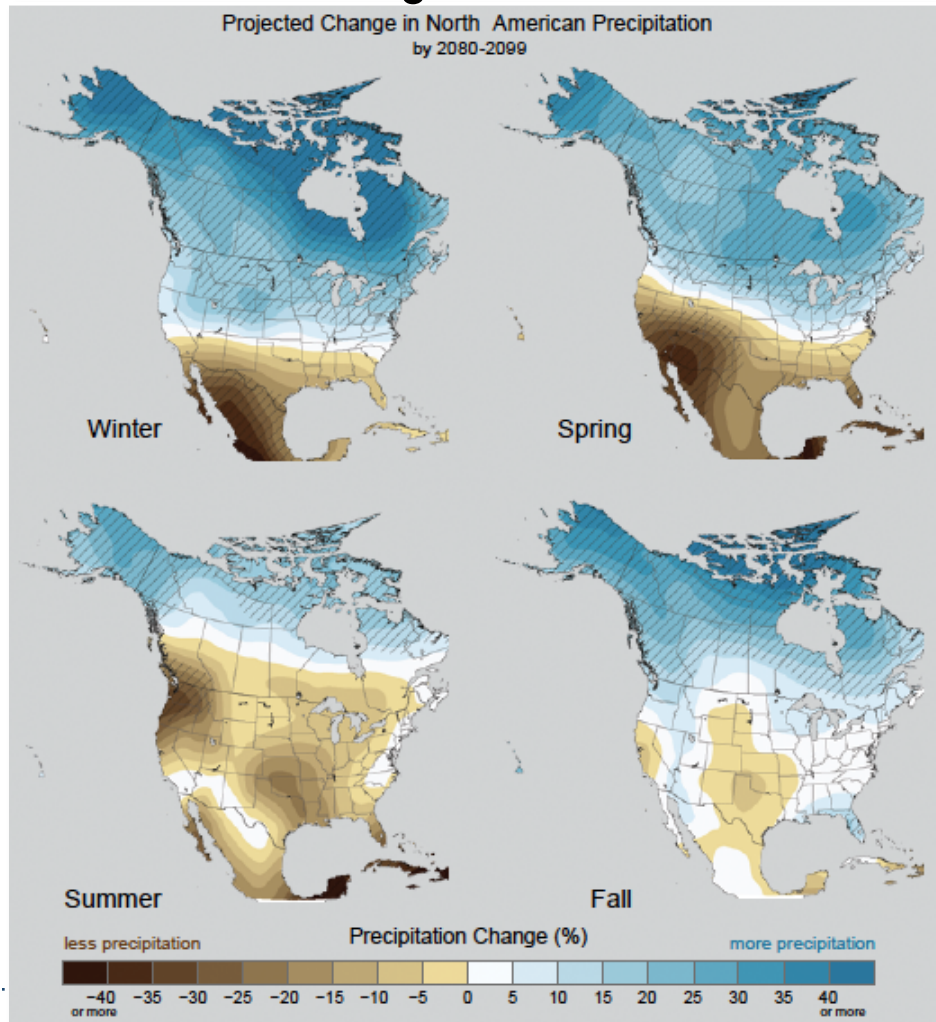


- 135 Tflops/s on 36,864 processors of the quad-core Cray XT4 Franklin at NERSC, 40% efficiency
- 224 Tflops/s on 163,840 processors of the BlueGene/P Intrepid at ALCF, 40% efficiency
- 442 Tflops/s on 147,456 processors of the Cray XT5 Jaguar at NCCS, 33% efficiency

..... For the largest physical system (36,000 atoms),
LS3DF is 1000 times faster than direct DFT codes


Other Efforts in Computing Sciences

Climate Modeling of Extreme Events



Wind Turbine Simulations

- Need for *higher fidelity predictions* in computational mechanics
- Develop new high-order accurate methods for problems with
 - complex real-world geometries
 - turbulent flows, waves, multiple scales
 - non-linear and fluid-structure interactions
- Our new schemes, algorithms, and solvers make high-order discontinuous Galerkin methods practical for realistic problems



Vertical wind turbine (velocity plot)


PI: Per-Olof Persson, UCB, LBNL

Uncertainty Analysis and Quantification

Design Intent: 66% (ASHRAE 90.1); Measured 44%



LEED ratings are based on **design** performance, not **actual** performance (EUI = End Use Intensity)



- Failure Modes Arise from **Detrimental Sub-system Interactions**
- Changes made to envelope to improve structural integrity diminished integrity of thermal envelope (“retainer wall as a fin”)
- Lack of visibility of equipment status/operation and uncertainty in loads (plug, occupancy, leaks), leading to excess energy use

The barrier is identifying critical (controlling) parameters for large scale dynamical systems (including physics and networks) and then using information to create and deploy and operate robust solutions

Figures from WBCSD, DOE (Chu), Tsinghua (Yi), NREL

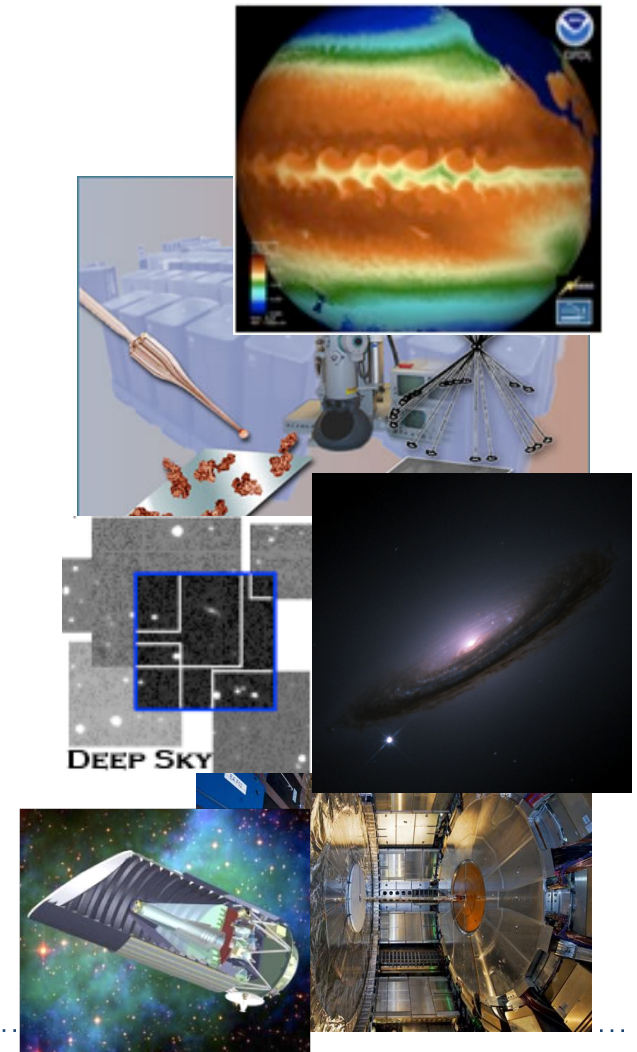
Computing for CC2.0| February 3, 2010

Prof. Yi Jiang, Tsinghua Univ. (2007)



Data Driven Science

- Ability to generate data is exceeding our ability to store and analyze it
 - Simulation systems and some observational devices grow in capability with Moore's Law
- Opportunity to lead creation of scientific communities around data sets
- A *science gateway* is a set of hardware and software for remote data/services
- Petabyte data sets will be common:
 - *Climate modeling*: IPCC will be 10s of petabytes
 - *Genome*: Genomes will double each year
 - *Particle physics*: LHC is projected to produce 16 petabytes of data per year



Carbon Cycle 2.0 Initiative



“Supercomputer modeling and simulation are changing the face of science and sharpening America’s competitive edge.”

Secretary Steven Chu



Thank You